

HP BladeSystem p-Class power solutions

technology brief



Abstract.....	3
Acronyms in text.....	3
Overview of HP BladeSystem p-Class power infrastructure.....	3
Power supplies.....	3
Power enclosure.....	3
Power distribution	4
Communication within the power infrastructure.....	4
Power distribution within enhanced server blade enclosures.....	4
Power distribution within standard server blade enclosures	6
Power supply redundancy	6
AC input redundancy	6
How server blades power up	7
Manual power on	7
Special considerations for users of facility DC power.....	7
Combining standard and enhanced server blade enclosures.....	7
Firmware requirements	7
Redundancy requirements	8
Combining single-phase and three-phase power enclosures	9
Recommendations.....	9
Use three-phase power	9
Design for redundancy	9
Understand when to use enhanced server blade enclosures.....	9
Use the Sizing Utility	10
Understand how workloads affect power requirements	11
Obtain real-time power information from iLO.....	12
Properly configure power zones.....	12
Maximum AC redundant configurations for HP BladeSystem server blades	13
ProLiant BL20p G3 and ProLiant BL25p server blades with mini bus bars.....	14
ProLiant BL30p and ProLiant BL35p server blades with mini bus bars.....	15
Scalable bus bars using 208V input.....	16

Scalable bus bars using 220V input	17
Summary	17
Appendix A. Power distribution options	18
For more information	20
Call to action	20

Abstract

The HP BladeSystem is an integral part of the HP infrastructure for the Adaptive Enterprise. One benefit of the HP BladeSystem is its modular infrastructure that customers can deploy in standard racks, along with existing legacy servers and storage devices.

With the introduction of the ProLiant BL30p (double-density) server blade, HP announced an enhanced server blade enclosure with a split power backplane. This technology brief describes the HP BladeSystem p-Class power subsystem and the different configurations that are possible when using both standard and enhanced server blade enclosures and different combinations of server blades.

This technology brief was written with the assumption that readers are familiar with the components of the HP BladeSystem p-Class infrastructure. For more information about the infrastructure components, see the HP website at www.hp.com/go/bladesystem/.

Acronyms in text

The following acronyms are used in the text of this document.

Table 1. Acronyms

Acronym	Acronym expansion
AC	Alternating current
DC	Direct current
EEPROM	Electrically erasable programmable read-only memory
RETMA	Radio Electronics Television Manufacturers Association (rack spacing)
VAC	Volts alternating current
VDC	Volts direct current

Overview of HP BladeSystem p-Class power infrastructure

The HP BladeSystem p-Class power infrastructure uses a rack-centralized power subsystem that consists of the power supplies, power enclosure, and the power distribution mechanism.

Power supplies

The capacity of each power supply varies with the voltage level of the local AC power source. Maximum capacity can only be achieved using a 240-volt nominally rated power source. Lower voltages may result in lower server blade capacity. Customers can use the [HP BladeSystem p-Class Sizing Utility](#)¹ to examine the effects of lower voltages. The “Recommendations” section in this paper describes the Sizing Utility.

Power enclosure

Customers can use either a single-phase power enclosure or a three-phase power enclosure. The power enclosures use two 30-amp, 200- to 240-VAC inputs.² The single-phase power enclosure holds up to four, 3000 watt hot-plug power supplies. The three-phase power enclosure holds up to six hot-

¹ The Sizing Utility is available at www.hp.com/products/servers/proliant-bl/p-class/info

² European systems use 32-amp, 200- to 240-VAC inputs.

plug power supplies. Therefore, HP recommends using the three-phase power enclosures because of its higher power capabilities, particularly to support the latest microprocessors or double-density server blades such as the ProLiant BL30p series.

Customers that have a direct facility power source of –48 VDC can use the DC Power Connection option kit to distribute power directly to the server blade enclosures without the use of the power enclosures. For more information, see the [HP website](#) or the “[HP BladeSystem p-Class System Overview and Planning](#)” document.

Power distribution

Power is carried from the power supplies in the power enclosure(s) to the server blade enclosures through one of three power distribution options: mini bus bars, scalable bus bars, or a bus box for a one-to-one connection. The bus bars are attached directly to the RETMA rails in a rack and provide power independently to each side (A and B) of the server blade enclosure. Appendix A gives additional details about the power distribution options.

Communication within the power infrastructure

Attached to the back of each server blade enclosure is a server blade management module. It is a self-contained microcontroller that communicates with the Integrated Lights-Out (iLO) management device on each server blade. The server blade management module also communicates with the power management module attached to the rear of the power supply enclosure. The power management module monitors the power supplies and power enclosure, delivers alerts and status data to the server blade management module, and determines during the auto power-up sequence if adequate power is available for newly installed server blades.

Power distribution within enhanced server blade enclosures

In 2004, HP introduced the HP BladeSystem p-Class enhanced server blade enclosure. This enhanced enclosure provides the power capacity to support the ProLiant BL30p and ProLiant BL35p series double-density server blades, well as supporting the higher power requirements of selected BL20p G3 and future HP server blades. Refer to the [HP BladeSystem enclosure compatibility matrix](#) for information about which server blade models are supported in the enhanced server blade enclosure.³

In addition, the enhanced server blade enclosure consolidates the iLO management ports through an internal unmanaged hub into a single iLO port. This port connects to all iLOs within the enclosure to reduce the number of management cables needed.

Two double-density server blades fit into a p-Class sleeve, which then slides into a single bay in the enhanced server blade enclosure. To support double-density server blades, the enhanced server blade enclosure has a signal and power backplane (referred to as the split power backplane) that allows for the physical connection of and support for the power loads of up to 16 server blades in a single server blade enclosure.

The enhanced server blade enclosure with the split power backplane is able to:

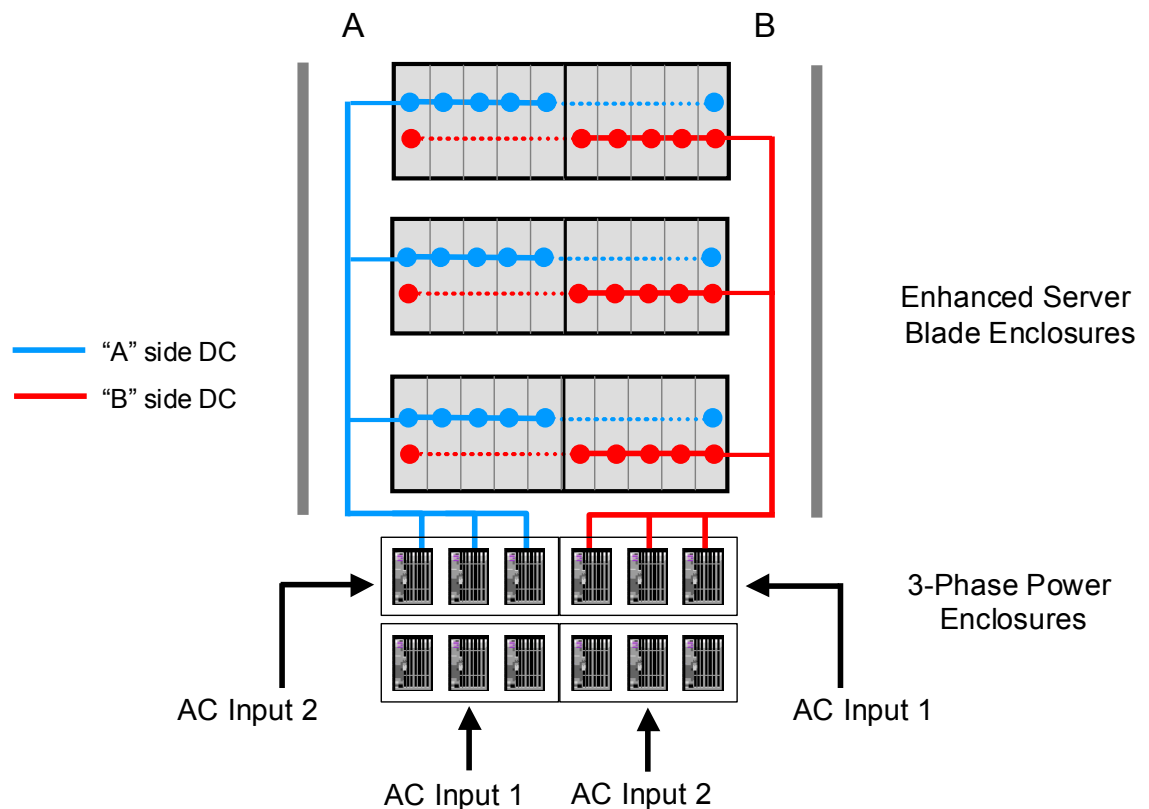
- Route power to up to 16 server blades.
- Support a higher input current and total input power than the shared power backplane in the standard server blade enclosure.
- Separate the DC power into the right and left sides of the server blade enclosure so that half the blades in a server blade enclosure are powered by side A (bays 1 – 4) and half by side B (bays 5 – 8). See Figure 1.

³ The compatibility matrix is available at <http://h18004.www1.hp.com/products/blades/components/Compatibility-Matrix.html>

It is important to note that the interconnect bays and the enclosure management module continue to share power from both sides (A and B) of the power subsystem. This allows both interconnect modules as well as internal chassis communication to continue in the event of a power failure from a single AC input.

Because the enhanced server blade enclosure splits the DC power into the left and right sides of each server blade enclosure, customers must use two power enclosures to configure AC line cord redundancy, as shown in Figure 1.

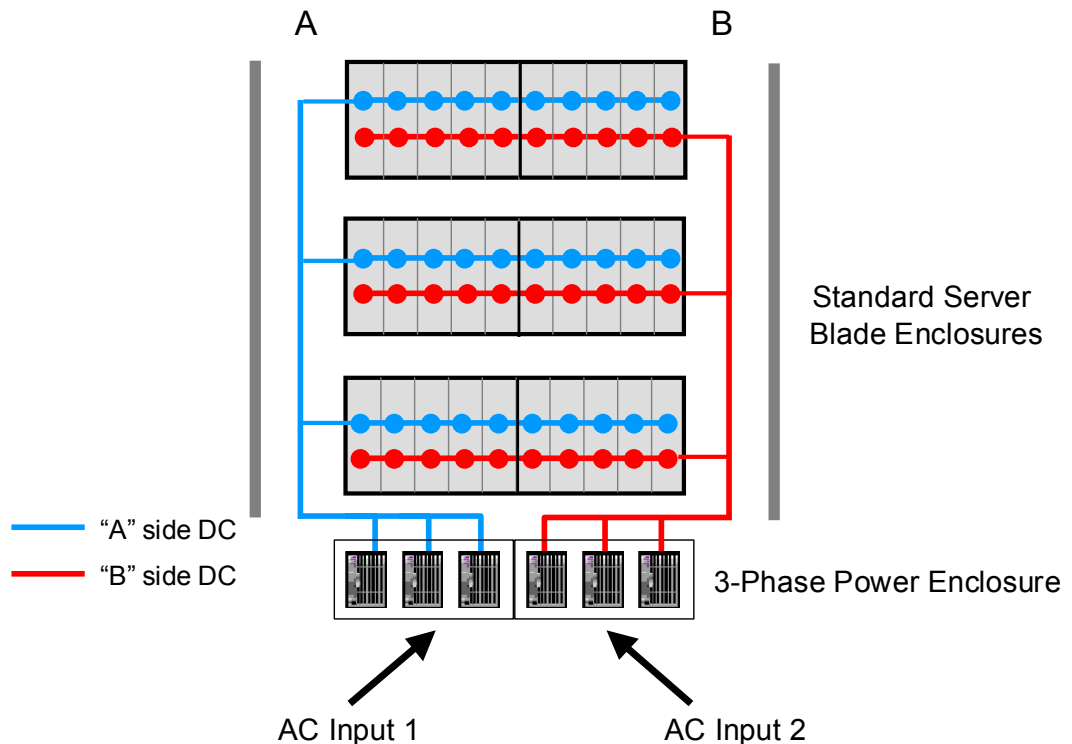
Figure 1. In the enhanced server blade enclosure, power supplies on side A distribute power only to side A server blades, while power supplies on side B distribute power only to side B server blades. Both sides A and B supply power to the interconnect bays and the enclosure management module.



Power distribution within standard server blade enclosures

Within the standard server blade enclosure, a shared power backplane distributes the DC power from each independent bus bar (or bus box) across all ten server blade and interconnect bays (Figure 2).

Figure 2. In the standard server blade enclosure(s), the shared power backplane distributes power from each side (A and B) of the power supply enclosure across all ten bays of the server blade enclosure.



The standard server blade enclosure supports first, second, and selected third generation ProLiant BL20p blades, as well as the ProLiant BL40p series server blades. Refer to the HP BladeSystem enclosure [compatibility matrix](#) for information on which server blade models are supported in the standard server blade enclosure.⁴

Power supply redundancy

Supporting power supply (DC) redundancy with standard server blade enclosures is possible when using a single power enclosure. Customers should use the [HP BladeSystem p-Class Sizing Utility](#)⁵ to determine power needs of specific configurations.

AC input redundancy

To ensure AC line cord redundancy when using standard server blade enclosures, each side of the power enclosure must be able to support the full power load of the blade enclosures. Therefore, the draw on each power feed must be limited to half its maximum power rating.

⁴ The compatibility matrix is available at <http://h18004.www1.hp.com/products/blades/components/Compatibility-Matrix.html>

⁵ The Sizing Utility is available from the website at: <http://www.hp.com/go/bladeSystem> by following the planning tools link.

How server blades power up

For a p-Class server blade to power on, the necessary power must be available from the rack-centralized power subsystem. The iLO management processor on each server blade verifies with the power management module that there is sufficient power before automatically powering on the server blade.

Once iLO receives a power-on request, it reads the EEPROM on the server blade to find the maximum blade wattage (generally in the range of 280 to 340 W or 755 W for a BL40p). The iLO management device sends a request for that wattage amount to the enclosure management module in the server blade enclosure, which then forwards the request to the power management module in the power enclosure. If—without exceeding the maximum power load—sufficient power is available to operate the additional server blade at its peak load and to meet redundancy requirements, then the power subsystem signals that the server blade may power up. If there is not enough power available, the iLO management device continues to retry the power-on request at 15-second intervals. After one minute, iLO continues retrying the power-on request in five minute intervals to reduce traffic on the communication bus.

After a server blade powers up, iLO collects data about the server blade's actual hardware configuration. The data allows the power management module to adjust the power allocation for that server blade to the actual maximum wattage required.⁶ This increases the remaining power available for server blades installed subsequently.

Manual power on

An administrator may occasionally want to power on a server blade even when the power management module indicates there is insufficient power—for example, when a non-redundant power configuration is adequate. The iLO management device allows administrators to perform a power override on the server blade in either of two ways: by pressing the server blade power button down for at least 5 seconds or by selecting the manual override button from the iLO virtual power button. The manual override must be used very carefully to avoid possible loss of service and data.

Special considerations for users of facility DC power

It is important to note that if a customer uses a facility DC power source, iLO assumes that there is a limitless supply of -48 VDC. Therefore, the iLO management device on the server blade does no calculations to ensure that adequate power is available to power up another blade; a newly installed server blade is always allowed to power up after it has been configured to use facility DC.

Combining standard and enhanced server blade enclosures

Some customers may choose to deploy both standard and enhanced server blade enclosures in the same rack to share the same power subsystem. Before doing so, administrators should be aware of firmware and redundancy requirements.

Firmware requirements

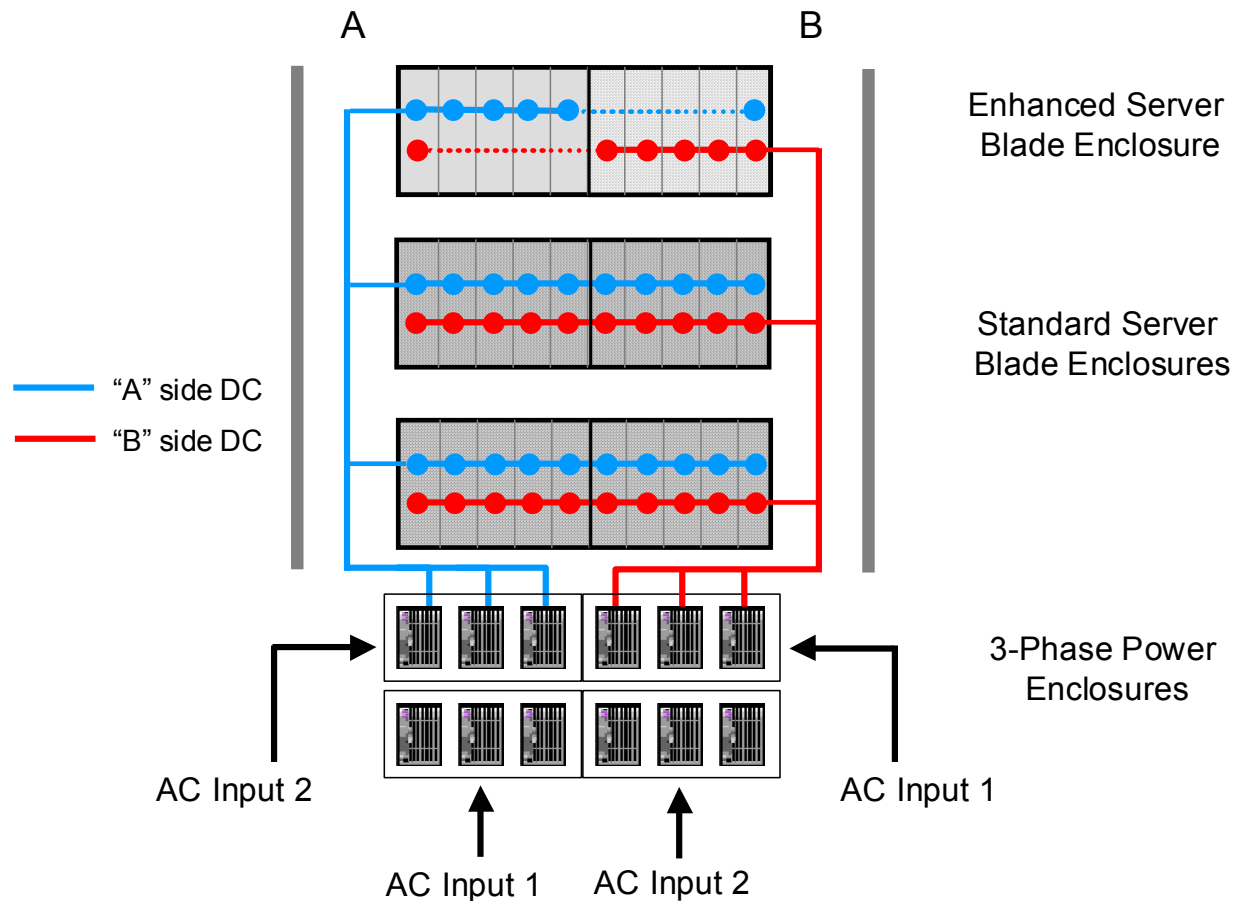
Enhanced server blade enclosures require firmware revision 2.03 or higher. The management modules on the server blade enclosure and on the power enclosure must have the same firmware version release.

⁶ This functionality is not available for the BL40p or the first generation BL20p.

Redundancy requirements

The power management firmware (version 2.03 or higher) calculates redundancy based on the type of server blade enclosure. For enhanced server blade enclosures, the power supplies on each side of the power enclosure are required to power the server blades on that same side, both the interconnect bays, and the enclosure management module. For standard server blade enclosures, each side of the power enclosure must be able to power up all the server blades and interconnects in that enclosure (see Figure 3).

Figure 3. Standard blade enclosures require both sides of the power enclosure to power up all the blades and interconnects in the standard enclosure.



Combining single-phase and three-phase power enclosures

It is possible to use both single-phase and three-phase power supply enclosures in a single power zone of a rack.⁷ If multiple power enclosures are combined in a rack in the same power zone, the power management firmware selects one of the power management modules to manage all the power redundancy calculations for that power zone. The power firmware selects which is the “master” power management module for a specific power zone. The master power management module keeps track of all power-related calculations for both power enclosures, such as amount of power available and redundancy requirements.

Recommendations

HP makes the following recommendations for configuring the HP BladeSystem p-Class power subsystem:

- Use three-phase power
- Design for redundancy
- Understand when to use enhanced server blade enclosures
- Use the HP BladeSystem p-Class Sizing Utility
- Understand how workloads affect power requirements
- Obtain real-time power information from iLO
- Configure power zones properly

Use three-phase power

HP BladeSystem p-Class systems require 30 or 32 amp, 200 to 240 VAC or a direct facility power source of –48 VDC power (± 10 percent). To gain the maximum benefit from the dense server blade designs, HP recommends using three-phase power to support the power loads.

Design for redundancy

To ensure continuous, redundant power, HP recommends using two power enclosures whenever a system includes enhanced server blade enclosures.

Understand when to use enhanced server blade enclosures

Customers must use the enhanced server blade enclosure when they deploy double-density blades or selected BL20p G3 and future HP server blades. Refer to the [HP BladeSystem enclosure compatibility matrix](#) for the most recent information about which server blade models require the enhanced server blade enclosure.

The standard server blade enclosure can be upgraded to an enhanced server blade enclosure by using the upgrade option kit. When upgrading standard server blade enclosures that use mini bus bars, the dual-power input kit for mini bus bars allows the mini bus bars to support a second power enclosure. This enables the power subsystem to provide maximum power redundancy using the enhanced server blade enclosures.

⁷ See the “Recommendations” section for additional details about the power zones.

Use the Sizing Utility

The [HP BladeSystem p-Class Sizing Utility](http://www.hp.com/go/ bladesystem/sizingutility) is a free tool available on the HP website.⁸ It provides valuable information for planning, ordering components, and preparing a site for delivery and installation of HP BladeSystem solutions. The user enters the desired server blade and enclosure configurations, chooses interconnects, and enters data center power information (Figure 4). Based on this user input, the Sizing Utility then calculates and displays:

- Maximum power specifications.
- Heat generation and cooling requirements
- Summary table of server blade components in the rack (server blades, memory, processor, etc.)
- Number of power supplies and power enclosures needed for both redundant and non-redundant power configurations
- System weight
- Equipment list

NOTE: The HP BladeSystem p-Class Sizing Utility assumes a maximum blade configuration and a 100 percent server utilization to calculate the wattage requirements and BTU output for a given rack configuration. This maximum load is rarely seen in typical operating environments. HP gives these maximum figures to help customers plan conservatively for their datacenter deployments.

⁸ The HP BladeSystem Sizing Utility is available at <http://www.hp.com/go/bladesystem/sizingutility>

Figure 4. Example of the HP BladeSystem p-Class Sizing Utility interface

Revision 0.46
February 2, 2005

HP BladeSystem p-Class Sizing Utility

Configurator Power Summary Equipment List Server Information Help

Enclosure 1 Enclosure 2 Enclosure 3 Enclosure 4 Enclosure 5 Enclosure 6 Rack & Power

Enclosure 1

RJ-45 Patch Panel 2

Make all Bays same as Bay 1

Server Blade Enclosure w/ Enhanced Backplanes w/ 8RDP licences

Clear Enclosure 1

Bay 1	Bay 2	Bay 3	Bay 4	Bay 5	Bay 6	Bay 7	Bay 8
TOP Blade	TOP Blade	TOP Blade	TOP Blade	TOP Blade	TOP Blade	TOP Blade	TOP Blade
BL35p	BL35p	BL35p	BL35p	BL35p	BL35p	BL35p	BL35p
374211-B21	374211-B21	374211-B21	374211-B21	374211-B21	374211-B21	374211-B21	374211-B21
1	1	1	1	1	1	1	1
AMD O250	AMD O250	AMD O250	AMD O250	AMD O250	AMD O250	AMD O250	AMD O250
2.4GHz-1MB	2.4GHz-1MB	2.4GHz-1MB	2.4GHz-1MB	2.4GHz-1MB	2.4GHz-1MB	2.4GHz-1MB	2.4GHz-1MB
0	0	0	0	0	0	0	0
None	None	None	None	None	None	None	None
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
Bottom Blade	Bottom Blade	Bottom Blade	Bottom Blade	Bottom Blade	Bottom Blade	Bottom Blade	Bottom Blade

Understand how workloads affect power requirements

Power requirements can be categorized into idle, maximum, and typical workloads.

Idle power requirements can be defined as the amount of wattage required to leave all servers powered on, but not executing any applications. In general, when all server blades connected to the power subsystem are idle, the HP BladeSystem power subsystem draws approximately 20 percent of the maximum power level.

The maximum load is calculated by assuming a maximum blade configuration and 100 percent average utilization of all the blades and enclosures connected to the power subsystem.

Typical, or actual, power draw is determined by the percent average server utilization of all applications running on the server blades. Therefore, typical power draws are unique to each customer's environment, and are usually much lower than the maximum power calculated by the Sizing Utility. Table 2 shows approximate percentages of the calculated maximum power based on average utilization rates. For this table, average utilization rate is defined as the average application utilization rate across all server blades connected to a specific power subsystem.

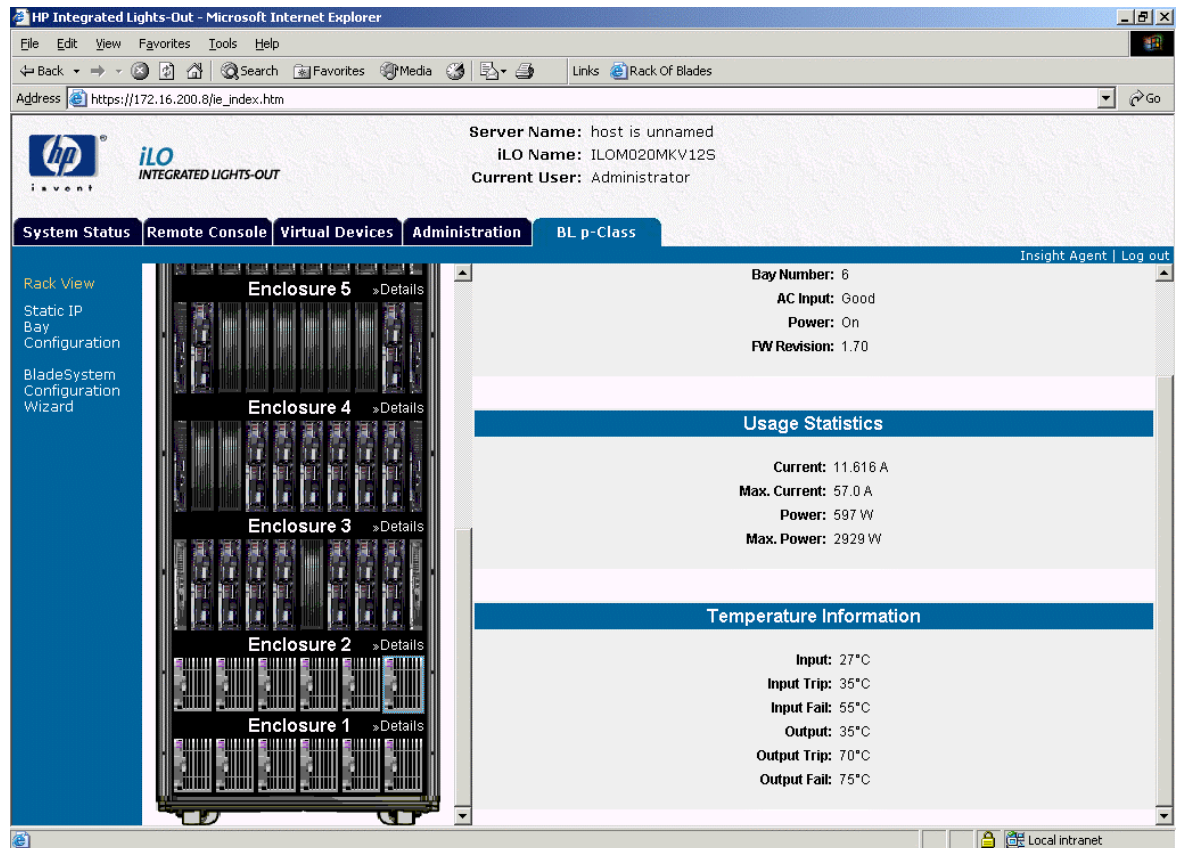
Table 2. General rules of thumb for correlating utilization rates to maximum power load

Average utilization rate	Approximate percent of maximum power load
Idle – 65% utilization	< 50% of max. power
65 – 80 % utilization	50 – 70% of max. power
99% utilization	90 % of max. power

Obtain real-time power information from iLO

Real-time power information about individual power supplies is available through the iLO management processor on each server blade. Figure 5 shows an example of the version 1.64 iLO interface.

Figure 5. Power information for the entire HP BladeSystem infrastructure is available through the iLO management processor on any installed server blade.




For diagnostic purposes only, power information is also available by using the serial port on the back of the power management module in each power enclosure. Users can connect a standard null modem cable to connect a laptop or other management computer to the serial port and access the information using a terminal emulator program. Data delivered through this port is not generally delivered to the management infrastructure.

Properly configure power zones

For a system to function properly when using two mini bus bar power subsystems in a rack, two separate power zones must be configured (Figure 6). Setting the power zone associates the server blade enclosures with the power enclosure(s) that support it. This allows power information and alerts to be sent to the proper power management module or server blade management module when power conditions change. For example, if an administrator removes a power supply from an enclosure, the information is sent to the server blade enclosures that are affected (the ones in that power zone).

Each power management module and each server blade management module has a two-position switch that sets the power zone for that module. All power configuration switches in the same zone must be set to the same position. Both standard and enhanced server blade enclosures use the power zone function in the same manner. For additional information about power zones, see the “[HP ProLiant BL p-Class System Maintenance and Service Guide](#)”.⁹

Figure 6. A fully populated 42U rack requires two separate power zones to identify which power enclosures supply specific server blade enclosures. This figure shows standard server blade enclosures; enhanced server blade enclosures use the same switches.



Item	Description
1	Zone 1 switches set in the down (default) position
2	Zone 2 switches set in the up (secondary) position

Maximum AC redundant configurations for HP BladeSystem server blades

Customers needing AC redundancy can use the following figures as general guidelines for maximum 42U and 21U rack configurations. The maximum configurations are based on either physical constraints or maximum power loads. The figures were calculated using these assumptions:

- All server blades in a rack are the same model
- All models are configured with the maximum number of processors, memory, and two disk drives

⁹ Available at <http://h200006.www2.hp.com/bc/docs/support/SupportManual/c00065042/c00065042.pdf>

- Therefore, actual server blade requirements may be significantly less than the maximum if server blades are configured with less hardware or used at lower average workloads. It follows that these maximum rack configurations can be exceeded if the blade servers are configured with fewer processors, less memory, and so on.

ProLiant BL20p G3 and ProLiant BL25p server blades with mini bus bars


Figure 7. Maximum configurations for ProLiant BL20p G3 or ProLiant BL25p server blades using three-phase power of 208V input or higher. Each configuration uses N+N redundant power supplies, redundant AC, and the mini-bus bars.



ProLiant BL30p and ProLiant BL35p server blades with mini bus bars

Figure 8 shows the maximum AC redundant configurations for either ProLiant BL30p or ProLiant BL35p server blades using the mini bus bars. The maximum configurations are the same for either 208V or 220V inputs.

Figure 8. Maximum configurations for ProLiant BL30p or ProLiant BL35p server blades using three-phase power of 208V input or higher. Each configuration uses N+N redundant power supplies, redundant AC, and the mini-bus bars.



Scalable bus bars using 208V input

Figure 9 shows the maximum AC redundant configurations when using the scalable bus bars and 208V input.

Figure 9. Maximum configurations for ProLiant BL p-Class server blades using scalable bus bars and three-phase, 208V power input. Each configuration uses N+N redundant power supplies and redundant AC.




Figure 10 shows the maximum AC redundant configurations when using the scalable bus bars and 220V input or higher.

42U

21U

server blade enclosure

8 8 8 8

8 8 8 8

40 blades

BL20p G3

8 8 8 8

8 8 8 8

40 blades

BL25p

12 16 16 16 16

60 blades

BL30p

12 16 16 16 16

76 blades

BL35p

The HP BladeSystem addresses customers' needs for increased density in overcrowded data centers. The enhanced server blade enclosure supports double density server blades as well as the latest generation models of other server blades. To supply power for such dense configurations, the enhanced server blade enclosure has a split power backplane that provides increased power capacity and splits the DC power across the two halves of the server blade enclosure. These design changes have important implications when deploying the HP BladeSystem server blades, such as configuring the power subsystem and configuring for redundancy. To ensure full power redundancy when using the enhanced server blade enclosures, HP recommends using two power enclosures and three-phase power in each configuration.

17

Appendix A. Power distribution options

The three options for HP BladeSystem p-Class power distribution are:

- Scalable bus bar
- Mini bus bar
- Power bus box

Table A-1 summarizes the characteristics of the power distribution solutions for HP BladeSystem p-Class server blades. Deploying an enhanced server blade enclosure requires two power enclosures to provide power redundancy.


Table A-1. HP BladeSystem p-Class power distribution options

Solution	Power enclosures supported	Server blade enclosures supported	Maximum rack space occupied
Scalable bus bar	2	5	36U
Mini bus bar	2 ¹	3	24U ²
Power bus box	1	1	9U

1. To attach two power enclosures to a mini bus bar, the Dual Power Input Kit for Mini Bus Bar option is required.

2. Deploying a full 42U rack of ProLiant BL p-Class server blades requires stacking two pairs of mini bus bars (Figure A-1).

Figure A-1. Scalable and mini bus bars



Item	Description
1	Scalable bus bars
2	Power supply enclosures
3	Dual Power Input Kit for Mini Bus Bar
4	Mini bus bars

For more information

For additional information, refer to the resources listed below.

Resource description	Web address
General HP BladeSystem information	http://www.hp.com/go/bladesystem/
"HP ProLiant BL p-Class System Overview and Planning" document	http://h200005.www2.hp.com/bc/docs/support/SupportManual/c00170396/c00170396.pdf Or by following the Technical documentation link on the general HP BladeSystem webpage
HP ProLiant BL p-Class Sizing Utility	www.hp.com/go/bladesystem/sizingutility/ Or by following the planning tools link on the general HP BladeSystem webpage
HP BladeSystem p-Class enclosure compatibility matrix	http://h18004.www1.hp.com/products/blades/components/Compatibility-Matrix.html
HP BladeSystem drivers and firmware	http://h71028.www7.hp.com/enterprise/cache/80365-0-0-0-121.aspx Follow this link, then click on the support & documentation link for the individual BladeSystem product
Technology briefs about HP BladeSystem	http://h18004.www1.hp.com/products/servers/technology/whitepapers/proliant-servers.html

Call to action

To help us better understand and meet your needs for ISS technology information, please send comments about this paper to: TechCom@HP.com